

AUTONOMOUS ONBOARD SCIENCE DATA ANALYSIS FOR COMET MISSIONS

David R. Thompson⁽¹⁾, Daniel Q. Tran⁽¹⁾, David McLaren⁽¹⁾, Steve A. Chien⁽¹⁾, Larry Bergman⁽¹⁾, Rebecca Castaño⁽¹⁾, Richard Doyle⁽¹⁾, Tara Estlin⁽¹⁾, Matthew Lenda⁽¹⁾,

(1) Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr. Pasadena, CA 91109, USA
Email: david.r.thompson@jpl.nasa.gov, all others *firstname.lastname@jpl.nasa.gov*

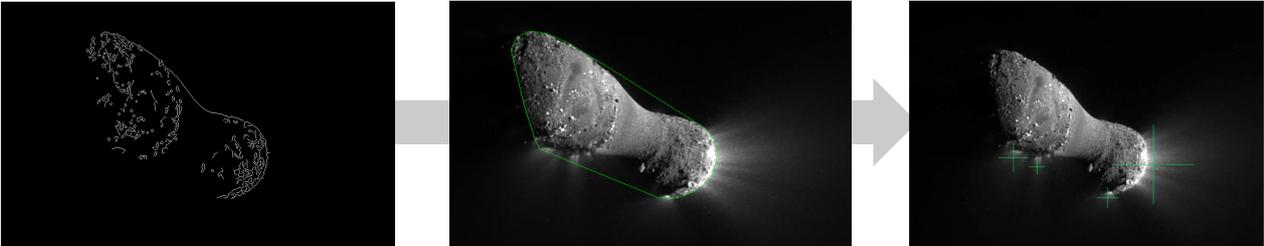


Figure 1. Plume detection by identifying the nucleus. Left: computational edge detection. Center: a convex hull of edge points. Right: bright areas outside the nucleus are plumes. Image credit: NASA/JPL/UMD.

ABSTRACT

Coming years will bring several comet rendezvous missions. The Rosetta spacecraft arrives at Comet 67P/Churyumov–Gerasimenko in 2014. Subsequent rendezvous might include a mission such as the proposed Comet Hopper with multiple surface landings, as well as Comet Nucleus Sample Return (CNSR) and Coma Rendezvous and Sample Return (CRSR). These encounters will begin to shed light on a population that, despite several previous flybys, remains mysterious and poorly understood. Scientists still have little direct knowledge of interactions between the nucleus and coma, their variation across different comets or their evolution over time. Activity may change on short timescales so it is challenging to characterize with scripted data acquisition. Here we investigate automatic onboard image analysis that could act faster than round-trip light time to capture unexpected outbursts and plume activity. We describe one edge-based method for detect comet nuclei and plumes, and test the approach on an existing catalog of comet images. Finally, we quantify benefits to specific measurement objectives by simulating a basic plume monitoring campaign.

1. BACKGROUND

Previous comet encounters include international flybys of 1P/Halley, the flyby of 81P/Wild by Stardust, the Deep Impact and NeXT encounters with 9P/Tempel 1, and a Deep Space 1 flyby of 19P/Borrelly. These few encounters have already revealed a very diverse population. Comets vary in size by orders of magnitude, with most having heterogeneous texture, albedo and composition. Each new visit reveals features not seen in previous cases. Tempel 1 has morphological evidence of

active geologic processes including scarps and outflows [1]. Its surface undergoes continuous modification, with visible change during the years between two flybys. The EPOXI flyby of comet Hartley 2 shows skyscraper-size spires, flat featureless plains that outgas H₂O, regions of rough and mottled texture, bands of various shapes, and diverse surface albedo. Comets' active areas range from 10-90%, changing over time and distance to the sun. They manifest as both localized jets and diffuse regions (Figure 1). Still more exotic, recently discovered “active asteroids” suggest that primitive ice could survive for billions of years in the inner solar system. This challenges the fundamental distinction between comets and asteroids [2].

A fundamental underlying question is the temporal evolution and driving mechanisms of comet surface activity. It is likely that comet activity changes on timescales faster than ground control's traditional command cycle and uplink interval. Faster reaction time will be important to characterize the dynamic activity profile. Additionally, future rendezvous missions would linger close to the surface where unexpected transient bursts could create navigation hazards for the spacecraft. Reacting on short timescales would thus be important to maintain safety during proximity operations.

Onboard data analysis can help address the challenge. Automated activity detection via morphology [3,4] or thermal properties [5] can analyze images onboard the spacecraft and prioritize active areas for selective caching and downlink. Instruments such as spectrometers or framing cameras can already collect more data than can be downlinked to Earth on tactical timescales [6]. Onboard detection could enable redundant image acquisition at high cadence, with

<p>Inputs: Set of image pixels $\mathbf{I} = \{i_1, i_1, \dots, i_n\}$; median filter width ψ; edge detection thresholds θ_1 and θ_2; plume detection threshold ϕ; a size threshold κ</p> <p>Outputs: A set of segments $\mathbf{S} = \{S_1, S_2, \dots, S_n\}$ corresponding to connected pixel regions containing plumes; a convex polygon comprised of pixels $\mathbf{P} = \{p_1, p_1, \dots, p_n\}$ containing the nucleus</p> <p>Begin DetectPlumes($\mathbf{I}, \psi, \theta_1, \theta_2, \phi, \kappa$)</p> <p>$\mathbf{F} \leftarrow \text{MedianFilter}(\mathbf{I}, \psi)$</p> <p>$\mathbf{E} = \{e : e \in \text{Canny}(\mathbf{F}, \theta_1, \theta_2)\}$</p> <p>$\mathbf{G} = \{g : g_i \in \text{ConvHull}(\mathbf{E})\}$</p> <p>$\mathbf{P} = \{p : p \in \mathbf{G}, p > \phi\}$</p> <p>$\mathbf{C} = \{C_i : C_i = \{c_i : c_i \in \mathbf{P}\}, c > \kappa, C_i \text{ is simply 4-connected}\}$</p> <p>$\mathbf{S} = \{s : s \in \mathbf{C}, s \text{ is not on nucleus or touching boundary}\}$</p> <p>Return \mathbf{S}, \mathbf{P}</p>	<p><i>compute median-filtered image</i></p> <p><i>compute the set of edge points</i></p> <p><i>compute convex hull of edge points</i></p> <p><i>image-wide intensity threshold</i></p> <p><i>connected components</i></p> <p><i>post-filtering rules</i></p>
---	--

Algorithm 2. Plume detection in a single image

priority transmission of any data containing outburst or plume events. Onboard detection would also permit adaptive instrument targeting. Here target selection in a context instrument directs a high-resolution sensor like a narrow- Field of View camera or spectrometer [7, 8,1]. For example, the Deep Impact Mission used a targeted slit observed the brightest image pixels to track a probe impact [1]. Active image analysis could extend this approach to track any transient outbursts. Finally, onboard processing could trigger evasion to rapidly avoid the most hazardous outbursts.

Onboard image analysis is a maturing technology. Previous studies used computer vision techniques to track asteroid surface features for navigation and odometry [9]. Planetary studies have also demonstrated reliable surface feature detection [7, 10] and crater detection [11,12]. Temporal change detection in remote sensing is broad topic with surface and deep space applications [13]. However, for comet applications the problem of separating plume from nucleus is fundamental to many science data collection activities. Moreover, it necessarily precedes many subsequent analysis tasks that focus on one or the other. The next section will detail one specific algorithm for detecting both plumes and nucleus.

2. PLUME DETECTION APPROACH

Several domain constraints guide our algorithm design. First, we desire that the method be computationally efficient for fast operation on radiation-certified spacecraft processors. It should provide robust performance across different illumination conditions. Spacecraft pose and the comet shape would both be uncertain during the initial characterization phase, so it is important that the algorithm not rely on strong assumptions about geometry. However, it is reasonable to expect that the image intensity might be normalized through *a priori* information or gain adjustment. Some limited parameter tuning for the specific characteristics

of the instrument should be possible in a real mission. However, we still desire robust performance for off-nominal imaging, new invented uses, or other violations of our base assumptions.

The proposed strategy finds the nucleus within each image and masks it; any bright remaining pixels correspond to plume activity (Figure) Our nucleus detection exploits the fact that comas are generally diffuse, while the nucleus is solid with a high-contrast terminator. Thus, *edge detection* reliably identifies points on the nucleus. After a preliminary denoising operation with a radius 5 median filter we apply the Canny edge detection algorithm [14]. This operation generally scales in proportion to image size. It uses some initial convolutional filtering operations and a subsequent hysteresis threshold that detects weak edges with support from their immediate neighbors.

The resulting set of edge points lies on the high-contrast horizon edge as well as interior edges from surface texture. We next form a nucleus “mask” by finding the subset of edge points that lie on the exterior contour. Planetary contours have smoothly-curved surfaces that can be found with a parametric fit [4] or Random Sample Consensus (RANSAC) approach [15,3] For comets with more irregular shapes, like the bilobate form of Hartley 2 (Figure), we model the nucleus using a conservative convex hull of edge points [16,3].

The 2D convex hull algorithm scales according to $O(n \log n)$ in the number of edge pixels, similar to an efficient sorting operation. This yields a subset of edge points whose polygon defines the nucleus mask. We apply an intensity threshold to the resulting exterior pixels and join them into connected components that constitute plumes (Figure right). As a final filter on poorly-localized plumes, we ignore any detection whose centroid lies on the nucleus or whose extent spans the entire image bounding box. Finally, we remove spurious detections by requiring each plume to have an area of at least κ (100 pixels in our case),

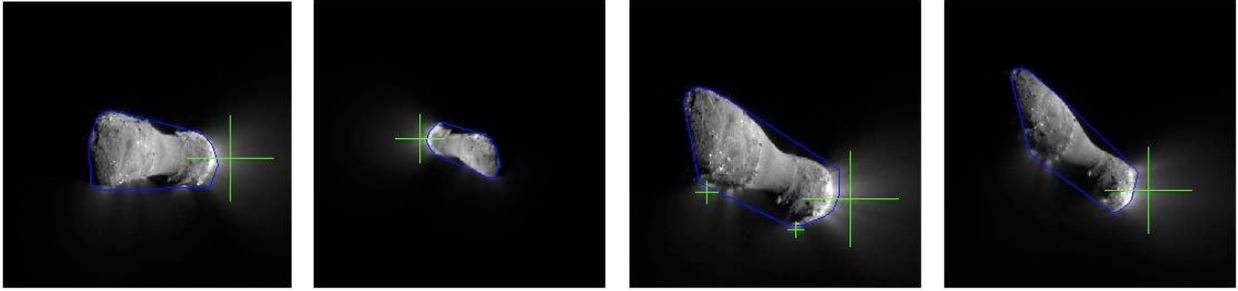


Figure 2. Comet plume detections from the EPOXI Hartley 2 sequence. Images from left to right: PIA13579, PIA13578, PIA 13570, PIA13600. Credit: NASA/JPL/UMD.

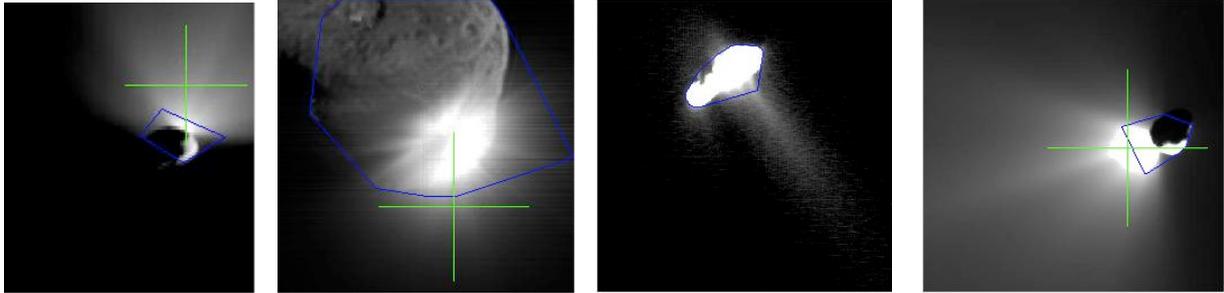


Figure 3. Comet images from Tempel 1 and Borrelley, Left to right: PIA02140, PIA202123, PIA03501, PIA02133. The saturated Borrelley plume fails the bounding box filtering rule and is not detected. Credit: NASA/JPL-Caltech/UMD.

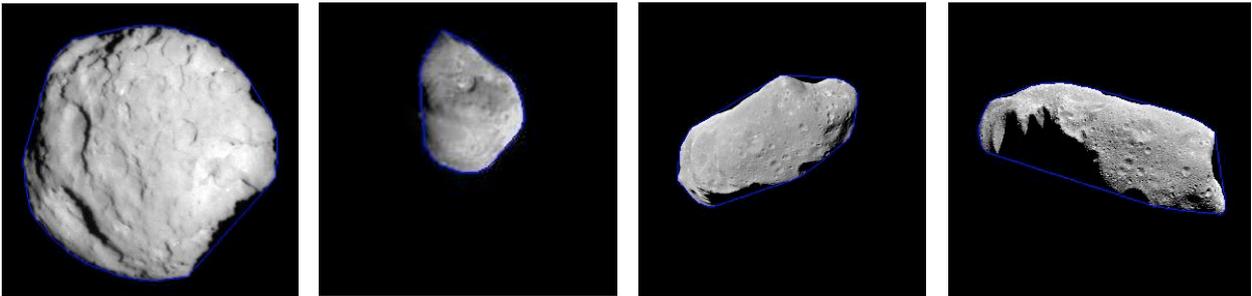


Figure 4: Several null cases with no obvious plumes. Left to right: PIA06285, PIA2125, PIA00136, PIA00135. Credit: NASA/JPL-Caltech.

Error! Reference source not found.Figure 1 illustrates the procedure. We refer the reader to Canny [14] for details of the “canny” edge detection subroutine, and Graham [16] for one convex hull algorithm suitable for the “ConvHull” subroutine. The important user-tunable parameters are the edge detection thresholds for the Canny procedure, the intensity threshold for detection and the size threshold κ . We find performance to be stable across a wide range of parameter thresholds.

Our tests evaluate plume detection performance on the existing catalog of cometary plume images, including images of Hartley 2 and Tempel 1. We developed the algorithm on image PIA13570 only, and then fixed these parameters for tests on the entire catalog. Figures 2, 3, and 4 show typical success and failure cases. Table 1 summarizes the results on the entire data set, and Figure 5 shows the computation time requirements as a function of image size. The plumes in the Tempel 1 images are the artificial result of the collision with the

Deep Impact impactor spacecraft, but we include these frames to demonstrate cross-scene performance. Failures occur only when the plume fills the majority of the frame.

3. SCIENCE YIELD

Here we aim to quantify the potential yield of an onboard activity monitoring campaign. We consider reference trajectories for the Rosetta encounter with Comet 67P/Churyumov–Gerasimenko in 2014. These represent a typical mapping scenario. The Rosetta operations plan calls for new trajectories to be specified at three-week intervals with more regular weekly updates to the data collection during each trajectory [17]. It is believed that there will be approximately three to five trajectory updates prior to the final delivery of the Philae lander to the surface. Plume activity bears on landing site safety so it will be particularly important to capture any image evidence of plumes during this

Image	Target	Plume present	Plume detection	$\sqrt{\text{(# pixels)}}$	Runtime (s)	Notes
PIA02124	Tempel 1			256	0.10	
PIA02125	Tempel 1			256	0.08	
PIA02140	Tempel 1	x	x	271	0.09	
PIA05571	Wild 2			370	0.21	a
PIA00228	Gaspra			400	0.15	
PIA02123	Tempel 1	x	x	471	0.32	
PIA03505	Borrelly	x		500	0.32	b
PIA03501	Borrelly	x		500	0.90	b
PIA03504	Borrelly			500	0.29	
PIA02127	Tempel 1			500	0.33	
PIA03500	Borrelly			500	0.29	
PIA13578	Hartley 2	x	x	501	0.32	
PIA13601	Hartley 2	x	x	501	0.23	
PIA13600	Hartley 2	x	x	501	0.28	
PIA13579	Hartley 2	x	x	501	0.31	
PIA13570	Hartley 2	x	x	501	0.28	c
PIA02133	Tempel 1	x	x	505	0.28	
PIA02134	Tempel 1	x		623	0.37	b
PIA00297	Dactyl			700	0.38	
PIA00069	Ida			769	0.37	
PIA02137	Tempel 1	x		900	0.91	b
PIA00299	Dactyl			1000	0.81	
PIA00118	Gaspra			1024	0.91	
PIA06285	Wild 2			1165	1.40	
PIA00136	Ida			1463	1.18	
PIA00135	Ida			2114	4.77	
PIA05004	Wild 2			2506	7.02	

Table 1 Performance. Notes: (a) diffuse coma, no obvious plume; (b) filtered by the bounding box rule; (c) training image used to develop the algorithm.

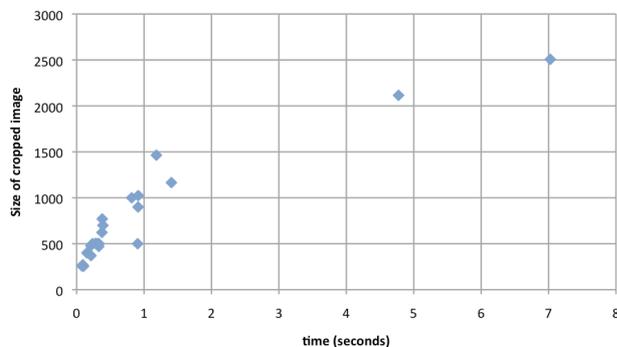


Figure 5. Time requirements vs. image size for the onboard plume detection algorithm, using the representative comet and asteroid images in Table 1.

initial period.

The spacecraft trajectories will be non-Keplerian, but a simple elliptical orbit segment can approximate a single trajectory's coverage. Figure 6 shows a hypothetical reference segment spanning the first three weeks of the characterization phase. This elliptical path encircles the nucleus along the terminator at an altitude of approximately 30km. The Rosetta spacecraft will follow each trajectory for three weeks at a time. Here we focus on the three-week interval from Nov. 15 to

Dec. 6 2014. Keeping to this interval, we can simulate the effect of a monitoring campaign while ignoring the more complex and mission-specific challenge of adaptive trajectory selection.

We model the nucleus a triaxial ellipsoid. Targets appear as points at randomly distributed latitude / longitude coordinates at an altitude of 200m above the surface (Figure 6 Right). We simulate 1000 events that appear at random times and last up to three weeks. We then apply geometric and illumination constraints to compute the time interval during which this target could be visible to the spacecraft. We perform these calculations with assistance from data and software provided by the NASA Navigation and Ancillary Information Facility [18]. We require that any detected plume be illuminated by the sun. It must not lie between the nucleus and the observer (in which case it is beneath the limb and therefore invisible to our approach). It cannot be occluded by the nucleus.

One key performance metric is the rate at which transient activity can be captured by followup data collection. Figure 7 shows the performance result of a monitoring approach in which a new jet triggers some dynamic reaction such as selective caching or followup. The left panel shows the fraction of captured targets that can be re-imaged during the same trajectory, showing a smooth decline due to occlusion and eventually a steeper degradation due to the finite lifespan of the event itself. The days of coverage possible for detected targets' followup measurements decline in proportion to reaction time. In addition to imaging the temporal evolution of the plume, it will be important for a monitoring campaign to capture many different spacecraft and solar angles to fully characterize plume dust reflectance properties. Figures 8 and 9 shows performance for these metrics. The wide range in science yield underscores the benefits of catching the plume as early as possible in order to improve the diversity of illumination conditions.

4. DISCUSSION

Overall these tests suggest that onboard activity detection could significantly improve science yield of a plume monitoring campaign. For activity events lasting 16 days, a hypothetical 14 day ground-in-the-loop planning cycle would permit active followup of approximately 25% of the actual events, while a 7-day planning cycle would increase this fraction to 65%. In contrast, onboard detection and followup would capture 95% of the events. The mean range of solar zenith angles for viewing each detected active region would increase from approximately 10 degrees (for the 14 day planning cycle), to 45 degrees per event. The contrast is even starker for shorter events, such as sub-day outbursts where onboard reaction appears vital to acquire any followup measurements at all.

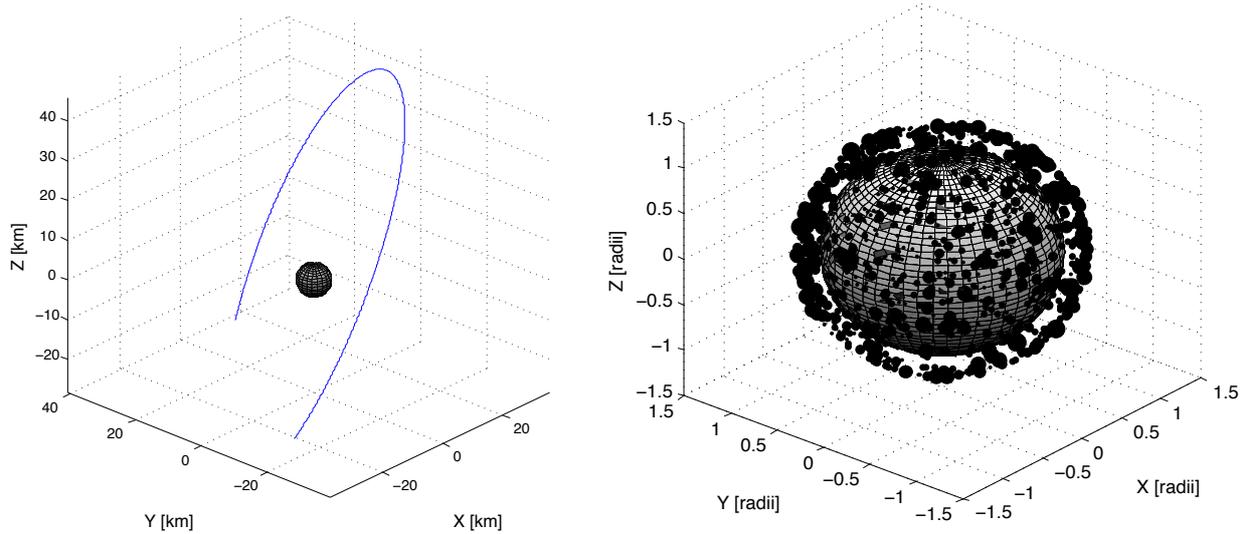


Figure 6: Simulated encounter. (Left) The trajectory of the spacecraft around the comet nucleus. (Right) Simulated surface activity events.

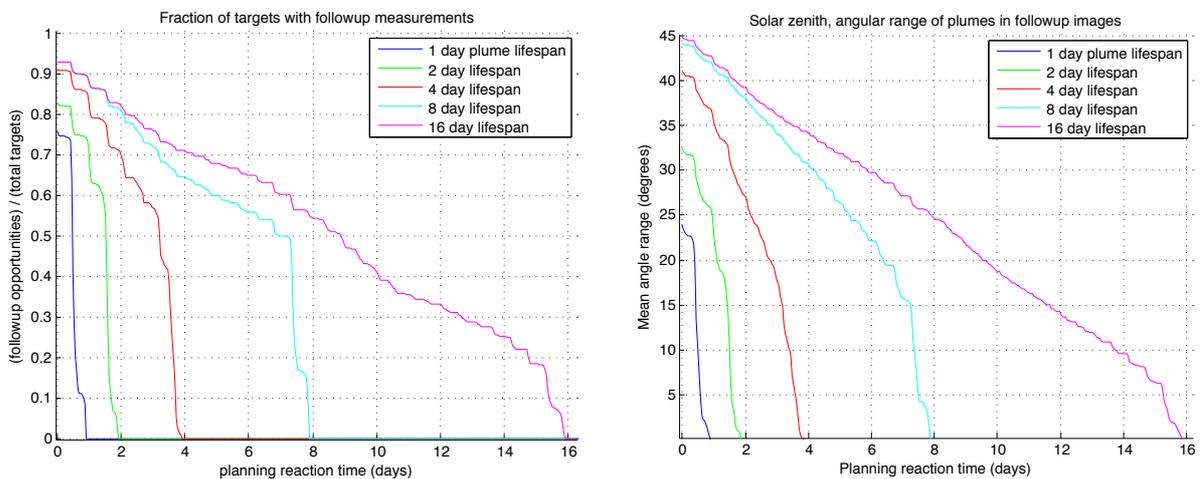


Figure 7: Performance vs. reaction time, for various automated plume monitoring strategies. Figures show the fraction of events captured in followup images (Left) and the span of solar zenith angles included in followup imagery (Right).

Naturally, other adaptive target selection approaches are possible and designers could complement morphological plume detection with temporal change detection or thermal analysis. Regardless of the specific approach, these simulations suggest onboard data analysis will be invaluable to understand transient activity on sub-day timescales. A monitoring campaign with a small optional resource allowance could considerably shorten the planning cycle, activating pre-validated followup options whenever plumes are detected. This could unlock the temporal dimension of cometary activity and significantly improve our understanding of the nucleus/coma interaction.

5. ACKNOWLEDGEMENTS

We thank the U.S. Rosetta team, Anna Marie Aguinaldo, Rob Witoff, Claudia Alexander, and Artur Chmielewski. We thank all the participants of the Agile Science study, including Tara Estlin and Kiri Wagstaff. This research was performed at the Jet Propulsion Laboratory, California Institute of Technology. We acknowledge U.S. Government Support. This study was conducted with sponsorship by the JPL Solar System Exploration Office and the JPL Research and Technology Development Program. Copyright 2012, California Institute of Technology. Government sponsorship acknowledged.

6. REFERENCES

1. A'Hearn M. F. et al., (2005) *Science* 310, 258 (10.1126/science.1118923).
2. Hsieh H. and Jewitt, D. The Search For Activity in 3200 Phaethon. *The Astrophysical Journal* (2005), 624. pp1093–1096.
3. Thompson, D. R., Bunte, M., Castano, R., Chien, S. and Greeley, R. (2012) Onboard Image Processing for Automatic Spacecraft Plume Detection, *Planetary and Space Science* 62. pp153-159.
4. Bue, B. Wagstaff, K. L. Castano, R. and Davies A (2007) Automatic Onboard Detection of Planetary Volcanism from Images. *Lunar and Planetary Sciences Conference*.
5. Castaño R., Wagstaff, K. L., Chien, S. A. Stough, T. M. and Tang, B. On-board Analysis of Uncalibrated Data for a Spacecraft at Mars. *KDD conference on Knowledge Discovery and Data Mining*, San Jose, California, 2007.
6. Gulick, V. C. Morris, R. L. Ruzon, M. A. Roush, T. L. (2001) Autonomous image analyses during the 1999 Marsokhod rover field test. *Journal of Geophysics Research*. 106E4. pp7745-7763.
7. Estlin, T. Bornstein, B. Gaines, D. Anderson, R. C. Thompson, D. R. Burl, M. Castano, R. and Judd, M., AEGIS Automated Targeting for the MER Opportunity Rover. *ACM Transactions on Intelligent Systems and Technology*, 2012 (to appear)
8. Thompson, D. R. Wettergreen, D. and Calderon P., F (2011). Autonomous Science for Large-Scale Robotic Survey. *Journal of Field Robotics*, Vol. 28, No. 4, July/Aug.
9. Udomkesmalee, S. Lin, C-F. Politopoulos, A.. Hu, G. Huntsberger . T. (2003) Autonomous Target Tracking for Asteroid Landing, Fourth International Conference on Control and Automation (ICCA-03).
10. Thompson D. R. and Castano B.. A Performance Comparison of Rock Detection Algorithms (2007). *Proceedings of the IEEE Aerospace Conference*.
11. Vinogradova T. Burl, M. Mjolsness, E. (2002) Training of a crater detection algorithm for Mars crater imagery. *Aerospace Conference Proceedings, IEEE*, vol.7, no., pp. 7-3201
12. Martins, R., Pina, P. Marques, J.S. Silveira, M. (2009) "Crater Detection by a Boosting Approach," *Geoscience and Remote Sensing Letters, IEEE*, vol.6, no.1, pp.127-131, Jan. 2009. doi: 10.1109/LGRS.2008.2006004
13. Castano, A. Fukunaga, A., Biesiadecki, J., Neakrase, L., Whelley, P. Greeley, R.. Lemmon, M. Castano, R. and Chien, S (2008) Automatic detection of dust devils and clouds on Mars, *Machine Vision and Applications* 19:5, 467-482, Springer.
14. Canny, J. (1986) A Computational Approach to Edge Detection. *Pattern Analysis and Machine Intelligence*, No. 6. pp 679-698.
15. Fischler, M. A. Bolles, R. C. (1981) Random sample consensus. *Communications of the ACM* 24 pp381–395.
16. Graham, R. L. (1972). An efficient algorithm for determining the convex hull of a planar set. *Information Processing Letters I*. 132- 133.
17. Koschny, D. Dhiri, V. Wirth, K. Zender, J. Solaz, R. Hoofs, R. Laureus, R.. Ho, T. M., Davidsson, B.. Schwehm. G. (2007) Scientific Planning and Commanding of the Rosetta Payload. *Space Science Reviews*. 128: 167–188.
18. Acton, C. H. et al., Ancillary data services of nasa's navigation and ancillary information facility (1996), *Planetary and Space Science* 44:1, pp. 65-70.